

ELUSIVE CAUSATION*

Cei Maslen

Victoria University of Wellington

cei.maslen@vuw.ac.nz

Talk at Sydney, July 19th, 2006

Intro

David Lewis claimed that knowledge is elusive. “That is how knowledge is elusive. Examine it, and straightway it vanishes...” He argued that epistemology robs us of our knowledge: “Maybe epistemology is the culprit. Maybe the extraordinary pastime robs us of our knowledge. Maybe we do know a lot in daily life; but maybe when we look hard at our knowledge it goes away.”

The aim of this paper is to answer the questions: Might causation be elusive in a similar sense to that in which knowledge has been claimed to be elusive? Might there be pastimes that rob us of causation too? A careful answer to these questions requires investigation into mechanisms by which the meaning of causal claims may be fixed by context. Hence, a sub-goal of this paper is to describe such mechanisms. Those familiar with Lewis’s paper “Elusive causation” will realize that the grunt work behind any proposal of a contextualist account is to describe such mechanisms and thus answer “*how can causation/knowledge/ability or whatever depend on context?*”

Section 1: Introducing Contextualism about Causation

In general, contextualism about a group of statements is the view that the meaning and truth conditions of statements in that group depend on facts about the situation in which those statements occur. For example, contextualism about knowledge is the view that statements of the form “S knows that P” and “S does not know that P” depend for their meaning and truth conditions on the context.

Contextualism about singular causal statements is the view that statements of the form “c is a cause of e” and “c is not a cause of e” depend for their truth-value and meaning on the context. The following examples illustrate the version of contextualism about causation I am interested in here. I begin with two examples of my own involving food, and follow with an example from Hartry Field involving an exploding gas tank.

Example 1: The Potato Eater

I was interested to read in a recent issue of *Scientific American* that potatoes are bad for you. The basic idea is that potatoes are wasted calories. They are bad for you, not because they directly have negative effects on your health, but because you could be eating something better instead.

To see how this illustrates contextualism about causation, let me describe a realistic conversation that might arise in reaction to this information. Suppose the potato-eater’s

* Thanks to John Carroll for letting me quote his unpublished work in this talk. Thanks also to my colleagues Tom Crisp and Josh Gert for patiently discussing some of these issues with me, and to Ian Rutherford for reading it through and encouraging me.

wife has read the *Scientific American* article and agrees with it. She says, “Yes, that’s right, eating potatoes every day is a cause of your bad health. If you’d listened to me and eaten broccoli instead, then you wouldn’t have had bad health.”

The potato-eater protests. “Of course I know that broccoli’s better for you than potatoes. But that doesn’t mean that potatoes are bad. If I hadn’t eaten potatoes, I wouldn’t have eaten broccoli instead. I’d have eaten something else. What if I’d eaten chocolate cake? Potatoes aren’t the cause of my bad health, ‘cause if I’d eaten chocolate cake instead of potatoes every day then I’d still have had bad health.”

It seems to me that both the potato-eater and his wife can be right, even though on the surface their claims appear to conflict. Eating potatoes every day, in contrast to eating broccoli every day is a cause of the potato-eater’s bad health. But eating potatoes every day, in contrast to eating chocolate cake every day, is *not* a cause of his bad health. The relevant alternative to the cause is determined by the context in which the causal claim occurs, and this has changed over the course of the conversation. For the potato-eater’s wife, broccoli is a relevant alternative to potatoes because she told her husband to eat broccoli instead of potatoes. For the potato-eater, eating chocolate cake is a relevant alternative to eating potatoes because that is what he (claims he) would have chosen to eat instead.

Example 2: Almond Crescents

Let’s suppose that I want to bake cookies, and Stuart has promised to try one. I only have two cookie recipes: almond crescents, and peanut butter dots. After much deliberation, I choose the almond crescent recipe, spend the day in the kitchen, and later Stuart comes over and tries one of my home-baked almond crescents. Unfortunately Stuart has a fatal nut allergy; he is allergic to all nuts. He eats the almond crescent and dies. Was my choice of the almond crescent recipe one of the causes of Stuart’s death?

Well, if I hadn’t chosen the almond crescent recipe then I would have chosen the peanut butter dot recipe because I wanted to bake cookies, and I only have two cookie recipes. So in one sense, choosing the almond crescent recipe was not a cause of Stuart’s death. If I hadn’t made almond crescents, I would have made peanut butter dots, Stuart would have tried one, and he would have died anyway. Relative to choosing the peanut butter dot recipe, choosing the almond crescent recipe was not a cause of his death. However, relative to the alternative of choosing to make no cookies, choosing the almond crescent recipe was a cause of his death. Taken out of context, there is no correct answer to the question whether my choice of the almond crescent recipe was a cause of Stuart’s death. Out of context, the only valid claims we can make are relative claims. However, the conversational context and unspoken assumptions can make some alternatives more salient than others.

Initially our awareness that I only have these two recipes may have made the peanut butter dot recipe a salient alternative for us; in this initial context the causal claim “my choice of the almond crescent recipe was a cause of Stuart’s death” is naturally interpreted as meaning “my choice of the almond crescent recipe in contrast to a choice of the peanut butter dot recipe was a cause of his illness.”

Example 3: The Exploding Gas Tank

The following example comes from Hartry Field.¹ A gas tank explodes on a jet. It is a moderately expensive model of gas tank, model T₂. Suppose that if the cheaper model T₁

¹ Fall 1997 seminar on causation at New York University.

had been used instead, it would still have exploded. And suppose that if the far more expensive model T_3 had been used instead, it would not have exploded.

Then it seems that in the context of a conversation in which the expensive model T_3 is considered a relevant option, the following sentence is true:

Using a model T_2 gas tank was a cause of the explosion.

However, in the context of a conversation in which only the cheaper model T_1 is considered a relevant option, the same sentence makes a false statement.

Note that there may be a correct answer to the question “If a model T_2 gas tank had not been used, which model of gas tank would have been used?” Suppose that the correct answer is T_1 : if a model T_2 gas tank had not been used then the cheaper model T_1 would have been used. But this alone doesn’t decide the question of whether using a model T_2 gas tank was a cause of the explosion. The contrast with using the T_3 model is still legitimate.

The main idea of the contextualist view I’m going to work with is that singular causal claims are implicitly contrastive in the way illustrated above. Singular causal claims are dependent on contrast events implicit in the context of occurrence of those claims. A popular approach to analyzing causation is in terms of counterfactuals. The basic counterfactual analysis says that one event is a cause of another if and only if had the first been absent the second would have been absent also. For example, the burning of the house was a cause of the roasting of the pig if and only if had the burning of the house not occurred the roasting of the pig would not have occurred either. This is clearly wrong for cases of preemption and overdetermination, but if you permit me to put these cases aside,² we can construct a simple contextualist view of causation by making a small amendment to the basic counterfactual analysis (on your handout):

“Event c is a cause of event e ” is true in context C iff had c not occurred and the *salient contrast to the cause* had occurred instead, then e would not have occurred and the *salient contrast to the effect* would have occurred instead.³

You might have been wondering how communication about causation can take place with so much left unsaid. After all, we are not mind readers. There is a lot to be said about this. Contrast events, if not communicated directly, may be clear to the audience by the previous course of the conversation, acknowledged assumptions, limitations, plans, and presuppositions. The set of contrast events is what Lewis calls a “component of conversational score”⁴. We have a tendency to interpret utterances generously or charitably, as true or probable, relevant, useful and informative. The set of contrast events is fixed and developed through the course of the conversation. In some cases the set is left unsettled or vague until a dispute arises.

² Let me be the first to admit that it may be a big mistake to put these cases aside here, since mental causation is one place where the threat of preemption and overdetermination has been posed. I have in the back of my mind that the correct treatment of preemption and overdetermination might slide into place by figuring out the mechanisms for fixing relevant contrasts that we employ. Fingers crossed.

³ More generally, ‘Event c , relative to a set of contrast events $\{c^*\}$, is a cause of event e , relative to a set of contrast events $\{e^*\}$, iff for every c^* in $\{c^*\}$ there is an e^* in $\{e^*\}$ such that if c^* had happened then e^* would have happened.’

⁴“Scorekeeping in a Language Game”.

You may have noticed that all three examples given above involve an agent's choice between a range of alternatives and then you may have wondered whether the context-dependence comes in at a different point - to determine the event of the choice. So in the first situation it could be that in one context the phrase "my choice of the almond crescent recipe" designates a contrastive event - my choice of the almond crescent recipe rather than the peanut dot recipe, and in another context the same phrase designates a different contrastive event – my choice of the almond crescent recipe rather than no recipe. Choices are naturally seen as contrastive - choosing the almond crescent recipe from one range of alternatives is naturally seen as a different act than choosing the almond crescent recipe from another range of alternatives. I don't have time to argue against this interpretation of the examples here. However, I think it is harder to interpret the context-dependence as determining a contrastive event in other examples where agent choices are not involved.⁵

A Different Version of Contextualism

I want to stress the difference between the contextualist view given here and an older contextualism concerning the notion of *the decisive cause*.

Mill had this to say on the process of selecting one causal factor from many as the decisive cause.

[T]hough we may think proper to give the name of cause to that one condition, the fulfillment of which completes the tale, and brings about the effect without further delay; this condition has really no closer relation to the effect than any of the other conditions has... Nothing can better show the absence of any scientific ground for the distinction between the cause of a phenomenon and its conditions, than the capricious manner in which we select from among the conditions that which we choose to denominate the cause.

[Mill, *A System of Logic*, ch. 5, §3]

Mill may have overstated the case. Selecting one causal factor from many as the decisive cause is an important part of historical interpretation; Mill is implying that historians make false or misleading causal claims almost all the time, and that there are not even rules to be found in their mistakes. Yet part of Mill's claim is clearly right and seems to have gained general acceptance. What counts as the decisive cause from among many causal factors is now generally recognized to be a heavily context dependent matter. Philosophical interest has turned almost exclusively to the liberal causal factor relation. But the examples in the previous section illustrate what is perhaps a more surprising fact: the notion of causal factor is itself context dependent. Whether something counts as a causal factor at all depends on features of the context.

I realize that you may not have been convinced by three examples that this simple contextual view is correct. However, the main purpose of this paper is not to argue for this view but to investigate its consequences, so let us press on.

⁵ Uriah Kriegel has argued for the view that context determines contrastive events in such cases. [personal communication]

Section 2: Mental Causation

You can probably already see how the truth of individual causal sentences is supposed to come and go on such a contextual view. An individual causal sentence may be true in one context and false in another. In order to understand whether causation is elusive in a similar grand manner to knowledge (as claimed), we must first determine details of the mechanisms for how relevant contrast events are fixed by a given context. We will focus in particular on the interesting case of mental causation. Let's also assume one position on mental causation: nonreductive materialism. That is the view that materialism is true and that mental properties are not reducible to physical properties.

Epiphenomenalism is the view that no (positive) claims of mental causation are true. Our question will be whether there are any contexts in which epiphenomenalism is true. This is the same as asking whether there are any contexts in which mental causation vanishes.

There is a point in most papers on mental causation where the author gestures at how many of our treasured beliefs are at risk of going down the drain when epiphenomenalism pulls out the plug. All I had written here was that it sounds wrong to say that my desire for a chocolate flake was not one of the causes of my buying of it. Jaegwon Kim is more inspiring, so let me just quote him: “[I]t is human knowledge and desire that built the pyramids of Egypt and the Great Wall of China; produced the glorious music, literature, and other artworks of our forebears; built our great cities; detonated nuclear bombs; and caused holes in the ozone layer.”⁶

There are several popular arguments for epiphenomenalism about mental properties: the Causal Exclusion Argument, the Inefficacy of Content argument, The Anomalism Argument, etc. But, let's leave for another occasion the question of whether contextualism about causation has a distinctive response to these arguments and apply the contextualist account directly to a few claims of mental causation.

First, a small complication. Nonreductionists often assume the following two-layered view of causation: *causation* holds between Davidsonian (coarse-grained) events, and a further causal relation, *causal relevance*, selects some properties of the cause and some of the effect as the efficacious ones. So corresponding to each cause-effect pair in the causal relation, there is a set of causally relevant properties of the cause and a set of causally relevant properties of the effect. Then even for nonreductionists who assume that mental events are identical to physical events, there is a threat of “property epiphenomenalism”, that is, the view that mental properties are not causally relevant.

The contextualist account of causation given above is just an account of event causation, but it could be expanded in various ways to add this extra layer of causal relevance. Instead I will try to sidestep this issue, and talk in terms of Kimian events, where a Kimian event is (roughly) an object's having of a property at a time. Nonreductionists should believe that Kimian mental events are distinct from Kimian physical events, so the epiphenomenalist worry is now that Kimian mental events are not causes.

In asking whether a mental event is a cause of a physical event, different contrasts may be relevant, depending on the conversational context. We may expect that in some contexts it will be true that a particular mental event causes a physical event; epiphenomenalism is false in those contexts. But perhaps there are some contexts in which no positive attributions of mental causation are true; epiphenomenalism will be true in such contexts.

⁶ From Kim, *Philosophy of Mind*, Westview Press, 1996, p. 125.

Consider my desire for chocolate. Was this one of the causes of my buying of chocolate? If instead of desiring chocolate, I had desired grapes, then instead of buying chocolate, I would have bought grapes. However, if instead of desiring chocolate, I was in the very same brain state, but it was totally lacking in any mental properties, then (presumably) I would still have bought chocolate. Or if instead of desiring chocolate, I was in the very same internal brain state, but external features of the world were the same as Twin Earth so that I desired shmocolate (made of XYZ), then (presumably) the same effect would have followed.

Here are some potential contrasts with my desire for chocolate:

- (1) Replacement mental event: The desire for grapes.
- (2) Replacement brain event: the firing of neuron₅₃₆ (instead of neuron₅₃₅)
- (3) Zombie event: An event physically identical to c, but lacking all mental properties
- (4) Ghost event: An event mentally identical to c, but lacking all physical properties
- (5) Twin-Earth replacement event: the desire for shmocolate
- (6) Range of Contrasts: a set of contrasts including all of the above
- (7) Zombie-world contrast: a contrast-world which is physically identical to the actual world but lacking any mental properties
- (8) Ghost-world contrast: a contrast-world which is mentally identical to the actual world but lacking any physical properties
- (9) Neuron-in-Petri-dish contrast: an event physically identical to c but placed in a very different extrinsic physical background⁷

But which contrasts are relevant, and in which contexts? Does contextualism about causation imply that epiphenomenalism is true in some contexts and false in others (that causation is elusive)? As I said before, the answer to this depends on the detailed mechanisms for determining relevant contrast events in a given context. There is a small industry devoted to finding such mechanisms for contextualism about knowledge. My hope in the following section is that recent work on the mechanisms of relevance for knowledge claims may carry over somewhat to the case of contextualism about causation.

Section 3: Impoverished Contexts

A first thought for how mental causation may vanish is that there are some contexts that are so impoverished as to fail to make salient any suitable contrasts with a mental event, and that all mental claims will be false in such contexts. (John Carroll suggested something like this in unpublished work, and this was my initial motivation for investigating this view.⁸) I will argue that this idea is incorrect but it will serve as a helpful starting point.

Consider the following two causal claims:

- (A) A desire for chocolate was a cause of her action.
- (B) The firing of neuron₅₃₅ was a cause of her action. (on your handout)

Perhaps there are contexts that are so impoverished as to fail to make salient *any* contrasts with the firing of neuron₅₃₅. Call these *neuroscientifically-impoverished contexts*. Similarly, perhaps

⁷ Karen Bennett discusses this alternative in her paper, "Why the exclusion problem seems intractable, and how, just maybe, to tract it", *Nous* 37:3, 2003, 471-497.

⁸ "Making exclusion matter less" [manuscript].

there are contexts that are so impoverished as to fail to make salient *any* contrasts with the desire for chocolate. Call these *mentally-impoverished contexts*. Then the suggestion is that, in neuroscientifically-impoverished contexts, no claims like (B) are true, and in mentally-impoverished contexts no claims like (A) are true. Roughly speaking, the claim is that microphysical causation “vanishes” in ordinary contexts, while mental causation “vanishes” in the context of a conversation between neuroscientists.⁹

If this suggestion is correct then the truth of a causal claim may depend on the state of ignorance/knowledge of the speaker, which may seem worrying. It is also interesting that if this suggestion is correct then causation disappears due to an opposite mechanism to the mechanism that contextualists about knowledge employ when they claim that knowledge disappears. According to contextualists about knowledge, when we move to a context with *more* relevant alternatives, knowledge vanishes; according to this suggestion, when we move to a context with *fewer* relevant contrasts, causation vanishes.

Although, as we shall see, this suggestion is mistaken, evaluating it carefully requires an understanding of the role of salience in fixing relevant contrasts. This is what we shall now turn to.

At first sight, it isn't very plausible that a conversation between neuroscientists is impoverished to the extent that no specific contrasts with the desire for chocolate are salient. Surely neuroscientists are equally aware as other human beings of alternatives to a desire for chocolate, and can easily call to mind some contrasting mental events such as the desire for grapes or the desire for a holiday.

However, on reflection, that depends on the way in which these contrasts are required to be salient. Although neuroscientists may *believe* that these other choices are available, they may not *have them in mind* or *be attending to them* while they're intensely discussing neuroscience. (One imagines here some absent-minded academics so steeped in intellectual discussion that they have difficulty locating the doorknob on their way out. This may not even be a foreign picture to us).

Here are some scenarios that determine various degrees and kinds of salience for proposition P for a speaker S.

- (a) P has never occurred to S, but S would agree to P if prompted. (We may say that S has the *dispositional belief* that P. For example, even those of us who have never read Fodor may have the dispositional belief that no grass grows on kangaroos.)
- (b) S is currently thinking about P.
- (c) Proposition P was recently mentioned in the conversation.
- (d) S is “seriously entertaining” P.
- (e) S is extremely anxious about P.¹⁰
- (f) S has regretted that not-P since the day he was born and his mission in life is to bring it about that P.

⁹ Let's distinguish between orders and levels. The threat of epiphenomenalism is usually presented as a threat against the causal efficacy of non-physical properties – the causal efficacy of all higher-level properties is supposed to be at risk of draining down to the level of physical properties. It is also sometimes presented as a threat against the causal efficacy of the macrophysical – the causal efficacy of the macrophysical is supposed to be at risk of draining down to the lowest microphysical level (if there is one.) Both macrophysical and higher-level properties are at risk if this story about impoverished contexts is correct. For more about the distinction between orders and levels see Huttenan and Papineau, “Physicalism decomposed”, *Analysis*, 2005.

¹⁰ Options (d) and (e) were inspired by Hawthorne's treatment of salience in *Knowledge and Lotteries*.

It may be true that, in the context of a conversation about neuroscience, no contrasts with “a desire for chocolate” are salient in the senses of (b) through (f). It may also be true that no specific contrasts with the firing of neuron₅₃₅ are salient in an ordinary context in any of the above senses. Specific neuroscientific properties do not come into mundane conversations, and ordinary people do not know enough about neuroscience to be having *secret* neuroscientific thoughts while carrying on mundane conversations.

However, there are still serious problems with this story. If no specific contrasts are salient in a context, then a causal claim doesn’t automatically become false or indeterminate. There are several reasons for this, and we will return to them later. Firstly, a reasonable rule of salience will make salience sufficient but not necessary for relevance. Salience will not be the only mechanism for making a contrast relevant in a context. Secondly, it seems to me that at least an *unspecific* contrast (“The absence of the firing of neuron₅₃₅”) is a relevant contrast to “The firing of neuron₅₃₅” in a neuroscientifically-impoverished context. Then with this unspecific contrast, statement (B) is true. Similarly, “the absence of a desire for chocolate” is a relevant contrast to “a desire for chocolate” in a mentally-impoverished context. Then with this unspecific contrast, statement (A) is true.

Section 4: Mechanisms for Establishing Relevance

Two Rules of Attention:

In the case of Contextualism about Knowledge, “The salience rule – or something like it – is the most common tool of contextualists.”¹¹ Lewis states the salience rule (or “rule of attention”) like this: “No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in this context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative.”¹² Other contextualists about knowledge dispute the details of the rule. Hawthorne argues that serious worrying is required. Just thinking about an alternative or mentioning it in passing doesn’t cut it. So, as Hawthorne says, merely watching *The Matrix* doesn’t rob one of the knowledge that one is at the movie theater. “Entertaining or attending to a state of affairs is one thing. Taking seriously the idea that things may be actually that way is quite another...”¹³

It also seems to me to be important to note that alternatives can be made salient in a way which makes it clear that they are *not* relevant. For example, someone finally says, “Forget BIVs for now, okay?” Or “Look, don’t talk to me about cleverly painted mules”. DeRose considers something like this; he talks about “veto power” and the “Aw, come on!” response to the mentioning of a skeptical hypothesis.¹⁴ So, we probably need two Rules of Attention for Contextualism about Knowledge:

- (1) Seriously entertaining that P is actually the case ensures that P is a relevant alternative.

¹¹ Hawthorne, *Knowledge and Lotteries*, p. 65.

¹² Lewis, “Elusive Knowledge”.

¹³ Hawthorne, *Knowledge and Lotteries*, p. 64.

¹⁴ “Single Scoreboard Semantics”.

(2) Seriously supposing that P is not actually the case ensures that P is *not* a relevant alternative.¹⁵

So what is a reasonable rule of salience for causation? Unfortunately we can't just copy rules straight from the contextualists about knowledge, because knowledge isn't causation. (Knowledge probably isn't even causation plus something else, to the disappointment of causal theorists of knowledge).

Also note that, whereas in the knowledge case to take an alternative as salient may be "to take seriously the idea that things may be actually that way", this cannot be so in the causation case. I know that I baked almond crescents and that I didn't bake peanut butter dots. That isn't in dispute. In the case of causation, to take an alternative as salient is to take seriously the idea that if the cause had not happened, things might have been that way. Or perhaps it is to take seriously the idea that these are the relevant ways that the cause might not have happened.¹⁶

Here, finally, is a suggestion for two Rules of Attention for Contextualism about Causation:

(1) Seriously supposing that "c is a relevant way for the cause to be absent" ensures that c is a relevant contrast.

(2) Seriously supposing that "c is *not* a relevant way for the cause to be absent" (e.g. c is not of interest) ensures that c is not a relevant contrast.

Consider the first rule. Hawthorne's insistence on serious entertaining or worrying seems reasonable in the causation case too. Merely mentioning a contrast out loud or having it flash through one's mind is not enough to make it a relevant contrast.

Consider the second rule. With Field's exploding gas tank example, this mechanism was employed to ensure that gas tank T3 was not relevant in one of the contexts considered. Field adverted to the great expense of tank T3. In one of the contexts there was presumably a tacit presupposition that only affordable alternatives are relevant. One imagines other practical reasons why a presupposition might be in place in a context that limits the relevant alternatives. There may also be explicit vetoes of alternatives. Perhaps S explicitly suggests that an alternative is relevant and T replies "Yeah, right. Like that was really going to happen".

¹⁵ Or perhaps, "A decision not to entertain that P is actually the case ensures that P is *not* a relevant alternative."

So what happens if both requirements are met? I don't know. Perhaps the "scoreboard explodes". DeRose suggests this for cases of interpersonal failure to agree on relevant alternatives. ["Single Scoreboard Semantics"].

¹⁶ This is another point where I disagree with Carroll. Carroll suggests the context-dependence of counterfactuals as a source of the context-dependence of causal claims. I agree, but I think that the context-dependence of causation goes a little further than this. He suggests fixing relevant contrasts by what is presupposed in each context as to what *would* have happened if the cause had been absent. But I think that there are contexts in which it is taken for granted that if Uncle Schlomo hadn't smoked two packs of cigarettes a day then he wouldn't have smoked at all, but in which the contrast with, say, smoking one pack a day is still a salient contrast. It is legitimate to contrast with any of a range of options that *might* have occurred if the cause had been absent. In the same vein, Carroll says "One thing the philosophers are likely to presuppose is that, if Harvey hadn't decided to raise his arm, then the underlying physical state would still have obtained". This does *not* seem to me likely to be a widespread presupposition amongst philosophers. Probably what is presupposed is that this is one relevant way for Harvey's decision to be absent.

The Second Rule of Attention is important. It will turn out that other mechanisms of relevance can make a range of alternatives relevant, not because they're being attended to, but for objective reasons, and then the Second Rule of Attention steps in to narrow this range.

Two More Rules

A Rule of Probability.

Context doesn't easily select far-fetched or improbable contrasts. This rule has both objective and subjective aspects. The probability of an alternative is an objective matter, but a threshold for how probable an alternative must be in order to count is a subjective matter that can differ between contexts.¹⁷

When I say (sarcastically) "Yeah, right. Like that was really going to happen" this seems to be both fixing a threshold for how probable an alternative must be in order to count and implying that it does not meet this threshold. This is a case where a Rule of Probability works together with an explicit veto works then.

Here is a suggestion for a Rule of Probability:

(3) An alternative that is very improbable, or below a threshold that is salient in the context, is not relevant.¹⁸

A Default Contrasts Rule

Maybe in the absence of other relevant contrasts one default *unspecific* absence of the cause is relevant.¹⁹ Or perhaps in that case a default *range* of contrasts is relevant: *all* the ways that it is physically possible that the cause could have been absent. I already mentioned this when discussing neuroscientifically-impooverished contexts. I pointed out that, in the absence of other salient contrasts at least an *unspecific* contrast ("The absence of the firing of neuron₅₃₅") is a relevant contrast to "The firing of neuron₅₃₅".

To say that a contrast, or a range of contrasts, is a default implies that this rule should come last in the pecking order. The Rule of Attention trumps all the other rules, and The Rule of Probability trumps the Default Contrasts Rule.

I'm not sure yet how the Default Contrasts Rule should go. Let's say...

(4) The unspecific contrast, the nonoccurrence of the cause, is relevant as a default.

Do we have to know what we're talking about?

¹⁷ Collins seems to have something like this in mind in his paper "Preemptive preemption" although he doesn't advertise his account as a contextualist account of causation. Also note that the Rule of Probability is similar to the main mechanism that Cohen postulates in his version of contextualism about knowledge. (See Hawthorne, *Knowledge and Lotteries*, p. 60, footnote 26.)

¹⁸ Note that the Rule of Probability is similar to the main mechanism that Cohen postulates in his version of contextualism about knowledge. (See Hawthorne, *Knowledge and Lotteries*, p. 60, footnote 26.)

¹⁹ Again the case of contextualism about knowledge provides an analogy. In Lewis's "Elusive Causation" he requires the alternatives to be highly specific. "If I ask you whether or not you have hands, I get you to entertain the proposition that you lack hands – but that does not, by Lewis's lights, make the possibility that you lack hands relevant." [Hawthorne, p. 63]. Of course, it would be fatal to the classic contextualist response to the skeptic if an unspecific contrast of this sort were always relevant by default.

It has been suggested to me that it doesn't seem right that the truth of a causal claim may depend on the state of ignorance/knowledge of the speaker.²⁰ This worry was introduced in section 3, and we are now returning to it.

Note first that causation doesn't really vanish in contexts in which the speaker is ignorant of available contrasts; it's just more difficult to convey certain things in these contexts. Just as knowledge doesn't really vanish in the context of a philosophy classroom; it's just more difficult to say what we want to say. (Rather, in different contexts, 2-place causal claims may pick out different 4-place causal relations.)

The Default Contrasts Rule is also helpful in responding to this worry. Even if the speaker doesn't know much about neurons, then it can still be true in her mouth that the firing of neuron₅₃₅ was a cause of some action. There will be default contrasts irregardless of the state of ignorance/knowledge of the speaker. Of course, when we don't know what we're talking about we may be unable to make salient some more esoteric contrasts. But our causal claims can still be interpreted straightforwardly and simply with default contrasts and may still be true with these contrasts.

Section Six: A Few Conclusions

It would be disappointing to end without applying these mechanisms for relevance to claims of mental causation. Let's concentrate on the two Rules of Attention, because they trump the other rules. Recall our examples:

- (A) A desire for chocolate was a cause of her action.
- (B) The firing of neuron₅₃₅ was a cause of her action.

Seriously supposing that any of our list of nine potential contrasts is a relevant way for the cause to be absent is enough to make it a relevant contrast. (One could do this by explicitly stating them with the right authority). So it is fairly easy to establish a context in which (A) or (B) is false. That is one way in which causation is elusive. And philosophers do often explicitly state relevant contrasts to mental events.

Consider a passage from Jaegwon Kim attempting to argue that mental properties are epiphenomenal in Davidson's picture:

[T]he very same network of causal relations would obtain in Davidson's world if you were to redistribute mental properties over its events any way you like; you would not disturb a single causal relation if you randomly and arbitrarily reassigned mental properties to events, *or even removed mentality entirely from the world.*²¹

One range of contrasts that this passage makes salient is the contrast with a range of worlds just like the actual world (or Davidson's version of the actual world) with mental properties redistributed over the events. With this contrast, epiphenomenalism is true. Another contrast that this passage makes salient is the contrast with a zombie world (a world with mentality entirely removed). With this contrast, epiphenomenalism is also true. Kim presumably doesn't believe that these are physically possible worlds, but by presenting them

²⁰ Maurice Goldsmith expressed this worry to me.

²¹ Kim, "The myth of nonreductive materialism", p. 269.

here he makes them salient contrasts nonetheless. So in the context established by this passage, epiphenomenalism is indeed true.

In conclusion, my discussion of the mechanisms of relevance shows that it's rather difficult to apply contextualism about causation to even one or two claims about mental causation. So it seems that we have a long way to go before we have the full story as to how causation is elusive.